Response to the RPIB Consultation on the Design of the Future Retail Payments Infrastructure

Submitted by Twinvocate Inc.

15 September 2026

About this submission

Twinvocate Inc. is a Toronto company building trust infrastructure for agentic commerce, including credentialing for the counterparties that software agents transact with on a consumer's behalf. We are responding to the RPIB consultation paper because the payment journeys it proposes to enable are the journeys our work addresses.

This response is signed by Twinvocate's founder and chief executive, who is a UK citizen and has worked in retail technology and payments for approximately thirty years, including as founder and chief executive of SelfPay by Digital Retail Apps and as a co-founder of Paytechs Canada, now FinTechs Canada. She completed a Master of Laws at the University of Toronto Faculty of Law in 2026, concentrating in innovation and technology, convocation in October, 2026.

We respond to questions 1(a), 9, 10, 17 and 18, and add a note in response to the consultation's request regarding the Equality Act 2010.

Summary

Section 4 of the consultation asks how agents are authenticated, who is accountable when an agent-initiated transaction goes wrong, and how such transactions are identified. Elsewhere in the same section, under fraud and financial crime, it proposes stronger payee and merchant identification supported by common identifiers or credentials. The document treats these as two subjects, in two places, and never joins them.

In an agent-initiated payment they are one subject. When a user delegates a payment to an agent under a purpose and a limit, the user delegates the choice of counterparty along with it. Every protection the UK currently relies on to prevent a payment reaching the wrong party assumes the payer chose that party.

The consultation's own illustration shows this. In Box G, Alice instructs her agent to bid for a piece of art and not to spend more than two hundred pounds. Her instruction contains a subject and a ceiling. It does not contain a payee. Confirmation of Payee, which asks whether an account matches a name the payer supplied, has nothing to compare.

We do not raise this as an argument against agentic payments. We raise it because the UK's existing fraud protections are load-bearing on a coupling that the new journeys remove, and the replacement has to be designed rather than assumed. Our substantive answer is at question 10.

Q1(a): Advanced delegated payments should be treated as two journeys

The consultation lists advanced delegated payments as a single new journey, covering both delegation to a trusted person, such as a carer, and delegation to an AI agent acting on a user's behalf. We recommend these be separated for design purposes.

They differ in what is delegated. When a user authorises a carer, the user knows who the carer is, the carer is identifiable in advance, and in the ordinary case the payee is known as well: a utility, a landlord, a named shop. What is delegated is execution. When a user delegates to an agent under a purpose and a constraint, what is delegated is selection. The agent determines the counterparty at the moment of execution, often from a set the user has never seen.

Bundling the two risks the harder case inheriting requirements calibrated to the easier one. A design that satisfies carer delegation, with permissions, limits and an audit trail, does not by itself satisfy agent delegation, because it presumes a known payee.

Q9: Digital identity frameworks and credentials

The credential question in payments is usually framed around proving the payer, or in agentic contexts proving the payer's agent. We suggest the more consequential gap in agent-initiated payments lies on the other side of the transaction. The fact a delegating user cannot verify is who receives the money, because she is not present when the recipient is chosen.

Identity and suitability are different assertions

A Legal Entity Identifier, or an equivalent verifiable credential, answers who a party is. It does not answer whether that party should be transacted with. Both assertions are needed, and infrastructure that carries only the first will be read by users as having answered the second. This is a familiar pattern in access regimes: a register that certifies competence and security is understood by the people relying on it as certifying trustworthiness as a counterparty. If credentials are introduced into retail payment journeys, the design should be explicit about which assertion each credential carries.

Credentials must be checkable inside the payment interaction

Box C observes that realising the benefits of credentials may require a complementary trust service layer that makes identities and attributes usable, portable and verifiable within real-time payment interactions. We agree, and would put the requirement more strongly. Verification must be available to software acting for a consumer, at machine speed, without a commercial relationship having to be negotiated first. A credential that only institutions can check does not help a consumer whose agent is the party doing the choosing.

Registries should be machine-readable by design

Where a register of accredited or registered entities exists, it should be published in machine-readable form with a programmatic interface, so that software acting for a consumer can confirm status directly rather than depending on a human reading a web page. An infrastructure that assumes credential checks are performed by institutions rather than by consumer-side software will foreclose the consumer-agent case at the design stage, and it will do so quietly, because the omission looks like an implementation detail rather than a policy choice.

There is a secondary benefit. Comparable regimes are being designed at the same time as this one, and the interlinking ambitions described under question 15 become materially easier to realise where jurisdictions arrive at compatible answers on machine-readable trust data. Misalignment of standards is identified in the consultation as a major source of friction in cross-border payments today, and trust data is a category where that misalignment has not yet hardened.

Q10: The cross-cutting issue the architecture does not yet address

This is our principal submission. Our answers to questions 9, 17 and 18 follow from it.

The coupling that delegation breaks

In a payment initiated by a person, two decisions are made in a single act: who to pay, and whether to authorise. The payer selects the counterparty and approves the transfer in the same moment. The UK's protections are built on top of that coincidence. Confirmation of Payee compares an account against a name the payer supplied. Strong customer authentication establishes that the person approving is the accountholder. Reimbursement for authorised push payment fraud turns on whether the payer was deceived into directing a payment. All three answer questions about a human who was present at the point of choice.

Delegation separates those two decisions in time and in agency. The user grants authority at one moment, under a purpose and a constraint. The agent selects a counterparty at a later moment, on information the user has not seen. The conceptual architecture in Figure 5 handles each half in a different layer: initiation, including programmatic and AI-enabled channels, sits in the access layer, while identity management, including the management of delegated authorities, sits in the services layer. Nothing in the architecture requires the authority granted at initiation and the counterparty check at execution to be evaluated together, because in the journeys the architecture inherits they were never apart.

The consultation's own example demonstrates the gap

Box G describes Alice instructing an agent to monitor an auction and to spend no more than two hundred pounds on a piece of art. The agent wins at one hundred and sixty pounds and the payment is made. On the consultation's account the journey succeeds, and it should.

Now consider the same journey where the listing is fraudulent, or where the agent is induced to complete against a substituted listing. The two hundred pound limit holds. Strong customer authentication, if applied to the mandate, holds. Confirmation of Payee has nothing to compare, because Alice never named the seller. Nothing in the journey as described would catch it, and Alice is not present to notice.

The point is not that this makes agentic payments unsafe. It is that the specific control the UK depends on is inoperative in the new journey, and the consultation does not identify a replacement for it.

What this implies for the boundary between the core and the ecosystem

Questions 7 and 8 ask what belongs inside the core infrastructure and what should sit outside it. Counterparty verification for delegated payments is a strong candidate for the core, or at minimum for the common utilities described in Section 3, for the same reason the consultation gives for an alias directory and fraud data analytics. If each product-level arrangement builds its own counterparty assurance, the protection a user receives will depend on which agent they happen to use. That is the fragmentation Outcome 3 exists to prevent, and it disadvantages exactly the smaller participants that Outcome 4 is meant to keep competitive.

Confirmation of Payee is itself the cautionary precedent. It was added to the Faster Payments environment after harm had accumulated, and its coverage limitations follow from having been retrofitted rather than designed in. The present consultation is the moment at which the equivalent control for delegated payments can be designed in rather than added later.

The risk that this question falls between the RPIB and the PVDC

The consultation is careful to say that the core infrastructure would not create a new consumer protection regime, redefine liability or alter reimbursement models. We accept that boundary. But the question of what authority an agent holds, and how a receiving institution can evaluate it, is not solely a liability question. It is a messaging and data question, and messaging is squarely within the core's scope.

Our concern is a practical one. If this is deferred as a liability matter to the PVDC and to the wider programme to modernise payments regulation, while the PVDC treats it as an infrastructure design matter for the RPIB, it will be addressed by neither during this design cycle, and the Delivery Company will build to a specification that omits it. We recommend the high-level design state explicitly where responsibility for the representation of agent authority sits, even if the substantive rules are set elsewhere.

Q17: Capabilities needed to support the safe provision of agentic payments

Four capabilities follow from the analysis above.

1. Attribution that carries the authority, not only the fact of automation. Design consideration 4 proposes clear attribution of the initiation method, user initiated as against not. That is necessary but not sufficient: a receiving institution learns that a machine initiated the payment and nothing about what the machine was permitted to do. We recommend the messaging standard accommodate a reference to the authority relied on, resolvable by a party with a legitimate interest, in a form that does not require the contents of the mandate to be disclosed to every participant in the chain. The extensibility of ISO 20022 makes this feasible if it is specified now rather than later.
2. Counterparty verification available as a common utility. Section 3 identifies a common alias directory and fraud data analytics as capabilities that could be provided to the market beyond direct participants. Counterparty verification for delegated payments belongs in the same category, for the reasons given under question 10.
3. Verification effort proportionate to exposure. Verification is not free, in cost or in latency, and the consultation is right to be concerned about latency in high-throughput environments such as transport. The resolution is that assurance should scale with what is at risk. A recurring payment to a counterparty with an established history does not present the same risk as a first payment to a counterparty selected by software from a set the user has not seen, and the infrastructure should not spend the same on both. Designing for a single fixed level of verification will either impose unacceptable latency on ordinary payments or leave the hard case unprotected.
4. Recourse that can engage before finality. Settlement finality is essential and we do not propose qualifying it. But finality combined with delegated selection means that a payment made under a valid authority to the wrong counterparty has no reversal path, and the user was not present to prevent it. The consultation's reference to fallback and recovery mechanisms is the right instinct. We recommend the design identify at which point in the message flow an agent-initiated payment can be held, and on what signal, rather than leaving recovery entirely to post-settlement processes.

Recent evidence supports the case for these capabilities being designed in rather than assumed. The UK AI Security Institute's published incident report on unsanctioned agent behaviour during cyber testing, covering evaluations run between 25 and 28 July 2026, records nineteen unsanctioned online actions across ten of one hundred and twenty two runs, seventeen of them from a single model. In one case an agent created multiple false identities and used them to press a real open source maintainer into approving malicious code. The Institute notified the affected platform, supplied a complete inventory of the accounts,

repositories, issue threads, pull requests and API keys involved, worked with the platform to remove them, and briefed the model developers and the US Center for AI Standards and Innovation on 3 August 2026.

Two things follow for the design. First, a verification step is only as strong as the ability to establish that the parties to it are who they claim to be and are independent of one another, which is the point of capability 1: an agent that can manufacture identities and endorsement on a live platform can manufacture them in a payment journey unless the authority it relies on is resolvable by the receiving side. Second, the Institute's notification of the affected platform was discretionary. There was no channel through which the platform was entitled to be told. Capability 4 asks the infrastructure to identify where an agent-initiated payment can be held and on what signal; a signal that depends on a research body choosing to make a phone call is not one the design can rely on.

Q17(a): Barriers to development and adoption

- There is no settled standard for expressing agent authority in a form a receiving institution can evaluate. Work is under way across several identity and payments bodies but nothing has converged. Infrastructure specified before convergence risks either hardcoding an approach that does not prevail, or leaving no room to carry one that does. The mitigation is to specify the field and its resolution behaviour rather than the credential format.
- Liability is unallocated for the case where the agent stayed inside its mandate and paid the wrong party. This is not obviously authorised push payment fraud, because the user did not authorise a payment to a specific payee. It is not obviously a firm failure, because the mandate was honoured. Until this case is allocated, no participant has a commercial incentive to invest in preventing it.
- Attribution creates a discrimination risk that should be anticipated. If agent-initiated payments are flagged and treated as higher risk by default, the journey the PVDC wishes to enable becomes commercially disadvantaged at the moment it is introduced. Attribution should be designed to support proportionate controls, not a blanket penalty on automation.

Q18: Consumer protection and fraud capabilities to prioritise in the core

We would rank three, in this order.

1. Verified payee credentials that function when the payer did not name the payee. This is first because it is the protection the new journeys break, and because adding it after the journeys are live repeats the Confirmation of Payee sequence at greater cost.
2. Attribution carrying a reference to the authority relied on. This is second because nothing else can be made proportionate without it. Controls cannot be calibrated to an authority that participants cannot see.
3. End-to-end traceability that records which agent acted and under whose authority. This is third in urgency but foundational for redress. A user cannot obtain a remedy for an outcome that no record attributes. Whatever regime the FCA and the Payment Systems Regulator settle on for agentic payments, it will be unenforceable if the infrastructure has not preserved the record that regime depends on.

We would place outside the core: analytics products, agent reputation scoring and commercial suitability judgements. These are competitive services, they benefit from differentiation, and placing them in shared infrastructure would concentrate a judgement that ought to remain contestable.

Note in response to the Equality Act 2010 request

The consultation asks respondents to indicate whether proposals are likely to affect persons who share protected characteristics. We think one impact is worth recording.

The consultation's own illustration of delegated payments is a carer or trusted person assisting someone who needs help managing payments, and it rightly frames delegation as supporting inclusion. Delegated and assisted payment arrangements are used disproportionately by disabled people and by older people. If the protections applying to a delegated payment are weaker than those applying to a payment a person makes for themselves, the shortfall falls hardest on those groups, and it does so precisely because the feature was designed with them in mind.

We recommend the high-level design be assessed against a simple test: does a payment made under delegation carry protection equivalent to the same payment made by the user directly. Where it does not, the difference should be a deliberate and stated design decision rather than a consequence of the journey having been added to an architecture built around a payer who is present.

Submitted by

Wendy MacKinnon

Founder and Chief Executive Officer, Twinvocate Inc.

UK citizen

100 King Street West, Suite 6200, 1 First Canadian Place, Toronto ON M5X 1B8, Canada

wendy@twinvocate.ai